# How Google's Networking Infrastructure Has Evolved Over The Last 10 Years

**Frederic Lardinois**
@fredericl / 3 years ago

**Google** today announced that it wants to fix your home Wi-Fi, but internally, it has long been working on far more complex networking issues. To connect the hundreds of thousands of machines that make up a Google data center, you can't just use a few basic routers and switches. So to manage all the data flowing between its servers, Google has been building its own hardware and software, and today, it's pulling back the curtain and giving us a look behind the evolution of its networking infrastructure.

Google's current setup, the so-called Jupiter network, has 100x more capacity than its first-gen network and can deliver a stunning 1 petabit/sec of total bisection bandwidth. The company says this kind of speed allows 100,000 servers to read all of the scanned data in the Library of Congress in less than a tenth of a second.

"Such network performance has been tremendously empowering for Google services," Google Fellow Amin Vahdat writes today. "Engineers were liberated from optimizing their code for various levels of bandwidth hierarchy. For example, initially there were painful tradeoffs with careful data locality and placement of servers connected to the same top of rack switch versus correlated failures caused by a single switch failure."

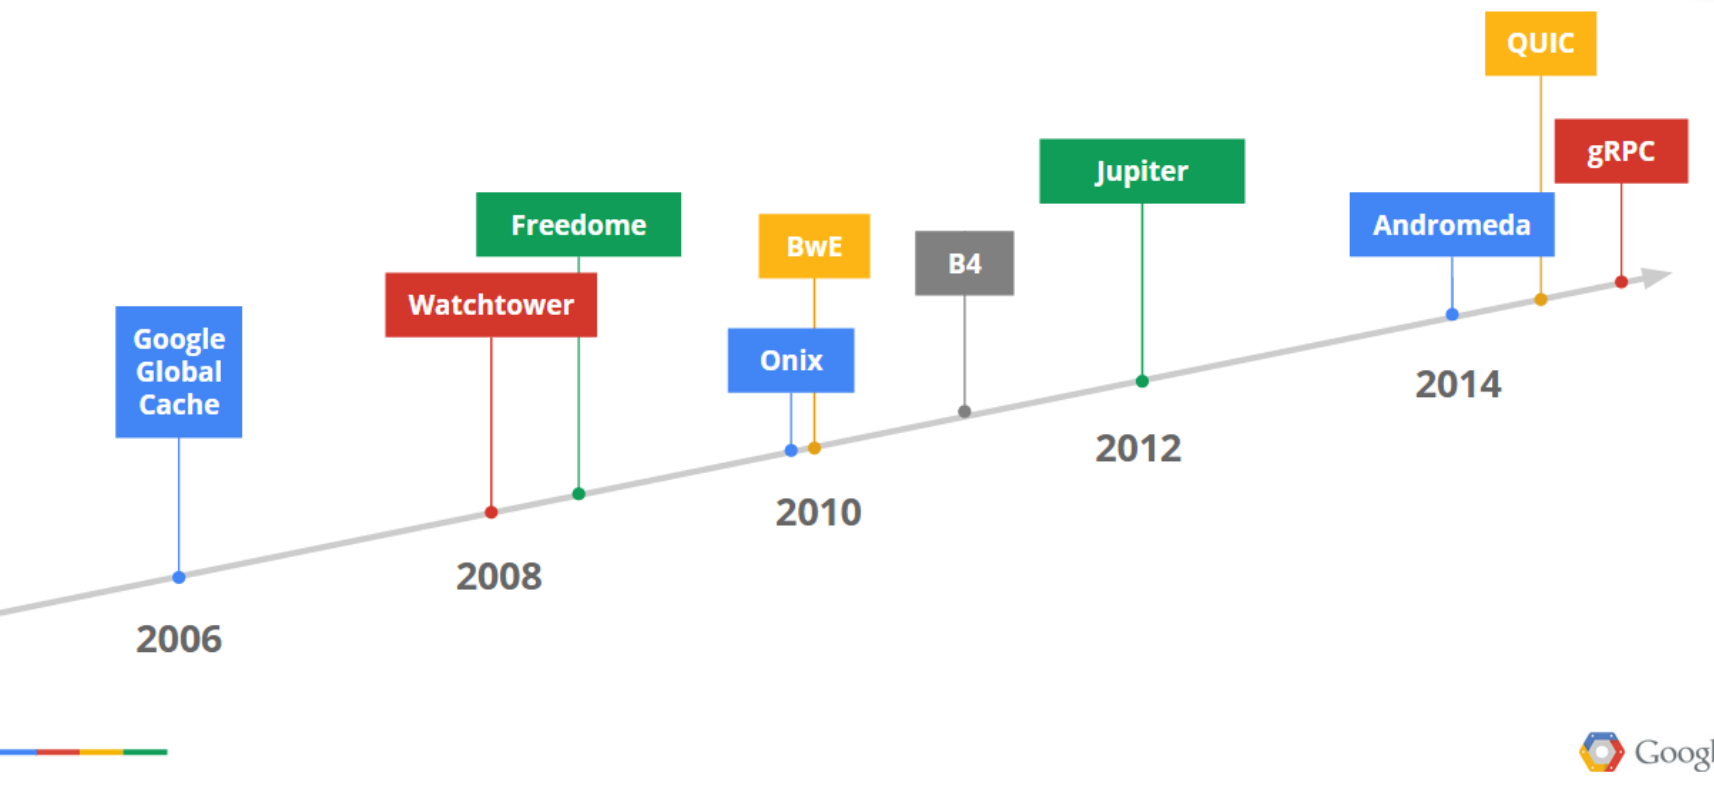


Ten years ago, however, Google didn't have anywhere near this capacity. This was before it had acquired **YouTube and** shortly after it launched products like **Gmail,** Google Earth and **Google Maps,** so the company's needs were likely changing very quickly.

Here is what those e **View 3 Photos** achines looked like:

As Google writes in a paper today, the company basically still deployed standard server clusters in 2004, and these 2005 machines are the first example of the networking it deployed in its Firehose 1.0 data center architecture. The goal of the 2005 machines was to achieve 1 Gbps of bisection bandwidth between 10,000 servers. To get there, Google tried to integrate the switching fabric right onto its homebuilt servers, but it turns out that "the uptime of servers was less than ideal."

With Firehose 1.1, Google then deployed its first custom data-center cluster fabric. "We had learned from FH1.0 to not use regular servers to house switch chips," the company's engineers write. Instead, Google built custom enclosures and moved to a so-called Clos architecture for its data center networks.

By 2008, Firehose 1.1 had evolved into WatchTower, which moved to using 10G fiber instead of traditional networking cables. Google rolled this version out to all of its data centers around the globe.

Here is what those racks looked like:

A year later, WatchT[View 3 Photos]norphed into Saturn. While the WatchTower fabric could scale to 87 Tb scaled up to 207 Tbps on denser racks.

Saturn clearly served Google well, because the company stuck with it for three years before its needs outstripped Saturn's capabilities.

"As bandwidth requirements per server continued to grow, so did the need for uniform bandwidth across all clusters in the data center. With the advent of dense 40G capable merchant silicon, we could consider expanding our Clos fabric across the entire data center subsuming the inter-cluster networking layer," Google's engineers write.

It's that kind of architecture that now allows Google to treat a single data center as one giant computer, with software handling the distribution of the compute and storage resources available on all of the different servers on the network.

The Jupiter hardware sure looks different from Google's earliest attempts at building its custom networking gear, but in many ways, it's also the fact that the company quickly adopted Software Defined Networking that allowed it to quickly innovate.

Google today released four papers that details various aspects of its networking setup. Because Google tends to hit the limits of traditional hardware and software architectures long before others, similar papers have often spawned a new round of innovation outside of the company, too.

It's unlikely that all that many startups need to build their own data centers, but other data center operators will surely study these papers in detail and maybe implement similar solutions over time (which in turn will benefit their users, too, of

course).

Add a Comment



## Sign up for Newsletters

See all newsletters

- [ ] The Daily Crunch
- [ ] The Weekly Roundup
- [ ] Crunchbase Daily

Email *

Subscribe

## Lime launches electric-assist bikes in its first UK city

51 mins ago **Catherine Shu**

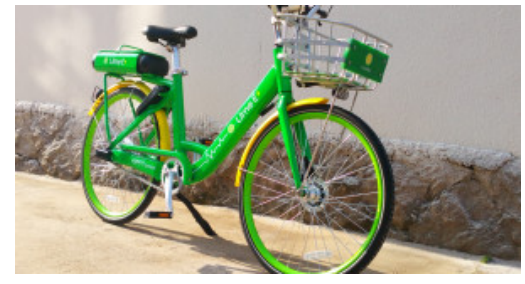

## WhatsApp's chief business officer is leaving

2 hours ago **Kate Clark**

## Enterprise messaging startup Eko Communications raises $20M for its European expansion

3 hours ago **Catherine Shu**

## Facebook denies report that election war room was disbanded

5 hours ago **Josh Constine**

## Apple launches app development program for female entrepreneurs

6 hours ago **Megan Rose Dickey**

## That night, a forest flew: DroneSeed is planting trees from the air

7 hours ago **Devin Coldewey**

## Apple to host free coding sessions at stores, rolls out new material for teachers

8 hours ago **Sarah Perez**

## Amazon launches a cloud-based robotics testing platform

9 hours ago **Brian Heater**

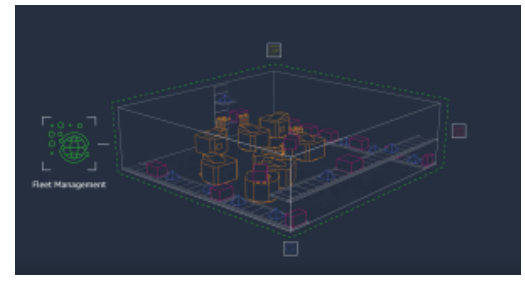

## GM plans to cut more than 14,000 jobs, close factories as downturn looms

10 hours ago **Kirsten Korosec**

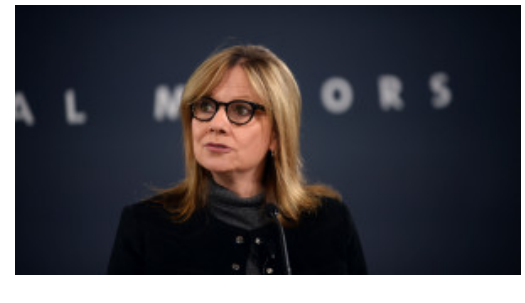

## Last-minute Cyber Monday deals

10 hours ago **Brian Heater**

## Google Maps biz reviews can now include hashtags

10 hours ago **Sarah Perez**


## Amazon says it's making freely available the same machine learning courses that it uses to teach its own engineers

10 hours ago **Connie Loizos**

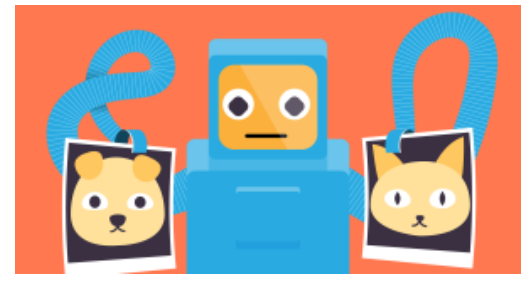

## Corvus Insurance lands a fresh $10 million to turn sensor data into actionable info for its food and pharma customers

11 hours ago **Connie Loizos**

## Watch NASA's InSight Mars lander touch down live (Update: Landed!)

11 hours ago **Jordan Crook**

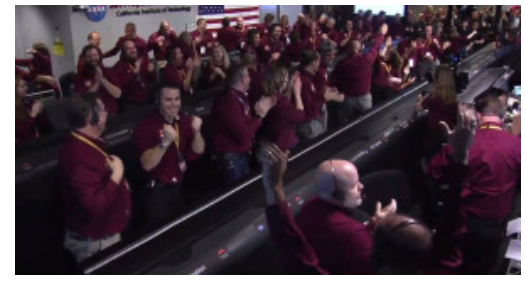

## Red Dead Online launches in beta tomorrow

11 hours ago **Brian Heater**

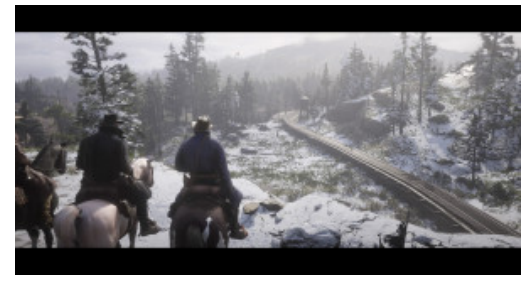

## Banuba raises $7M to supercharge any app or device with the ability to really see you

12 hours ago **Mike Butcher**

## Amazon closes its restaurant delivery service in London

12 hours ago **Sarah Perez**


## Coinbase and others invest $12.75M in project to create 'digital security offerings'

12 hours ago **Jon Russell**

## Cyber Monday projected to hit $7.8B in online sales, $2B spent via smartphones

12 hours ago **Ingrid Lunden**


## Amazon exec Dave Stephenson tapped to lead Airbnb to the public markets

13 hours ago **Kate Clark**